

Bounds on the Efficiency of Black-Box Commitment Schemes

Omer Horvitz¹

Jonathan Katz²

Department of Computer Science, University of Maryland, College Park, MD 20742, USA

Abstract

Constructions of cryptographic primitives based on general assumptions (e.g., one-way functions) tend to be less efficient than constructions based on specific (e.g., number-theoretic) assumptions. This has prompted a recent line of research aimed at investigating the best possible efficiency of (black-box) cryptographic constructions based on general assumptions. Here, we present bounds on the efficiency of statistically-binding commitment schemes constructed using black-box access to one-way permutations; our bounds are tight for the case of *perfectly*-binding schemes. Our bounds hold in an extension of the Impagliazzo-Rudich model: we show that any construction beating our bounds would imply the unconditional existence of a one-way function (from which a statistically-binding commitment scheme could be constructed “from scratch”).

Key words: Cryptography, commitment schemes

1. Introduction

A central focus of modern cryptography has been to identify the minimal assumptions needed for the construction of various cryptographic tools and protocols. We now know, for example, that one-way functions are necessary [IL89, Rom90] and sufficient for the construction of pseudorandom generators (PRGs) [BM84, Yao82, GL89, HILL99], universal one-way hash functions (UOWHFs) and digital signature schemes [NY89, Rom90], private-key encryption schemes [GGM85], and commitment schemes [Naor91]. Unfortunately, all the constructions just referenced are notoriously inefficient, and no constructions (based on one-way functions) improving upon the efficiency of these solutions

Email addresses: jkatz@cs.umd.edu (Jonathan Katz)

¹Work done while in the Department of Computer Science, University of Maryland. Research supported by U.S. Army Research Office award DAAD19-01-1-0494.

²Corresponding author. Research supported by NSF CAREER award #0447075.

are known. On the other hand, more efficient constructions *are* known to exist under stronger (e.g., number-theoretic) assumptions.

The apparent tradeoff between the efficiency of a construction and the underlying hardness assumption used to prove it secure has prompted a recent line of research aimed at answering the following question: *how efficient can constructions based on minimal assumptions be?* One way to formalize this question is to look at so-called “black-box” constructions which use an underlying primitive (e.g., a one-way permutation) only as an oracle, but do not require an explicit circuit computing the primitive in question (see Section 1.1 for further discussion). The idea of studying cryptographic constructions in this way was initiated by Impagliazzo and Rudich [IR89, Rud88] for proving *impossibility* of certain constructions, and much additional work in this vein followed [Rud91, Sim98, G⁺00, GMR01, Fis02]. (See [RTV04] for rigorous formal definitions of the Impagliazzo-Rudich model, as well as some variants that have been used.) Kim, Simon, and Tetali [KST99] were the first to use this model as a means of studying the *efficiency* of constructions (rather than their *feasibility*), with efficiency measured in terms of the number of oracle calls made by the construction. They showed non-tight bounds on the efficiency of constructing UOWHFs from one-way permutations. Extending their results, Gennaro, et al. [GGKT05] show that known constructions of UOWHFs based on one-way permutations are in fact optimal; they also show efficiency bounds for the case of PRGs, private-key encryption schemes, and digital signatures based on one-way permutations, as well as for the case of public-key encryption schemes based on trapdoor permutations.

In this work, we bound the efficiency of constructions of statistically-binding commitment schemes based on one-way permutations. Our results are tight (i.e., they match known constructions) for the case of *perfectly*-binding commitment schemes. We remark that such bounds do not follow from the work of Gennaro, et al. [GGKT05], and in fact proving bounds for the case of commitment schemes was left as an explicit open question there. Indeed, beyond the additional technical ideas used in this work, our bound is interesting as the first example of an efficiency bound on a protocol which protects against malicious participants (the cryptographic primitives considered in [KST99, GGKT05] only involve honest participants, with the adversary being a “passive observer”).

Before describing our results in more detail, we provide a brief overview of the Impagliazzo-Rudich model and black-box lower bounds. (The following is adapted from [GGKT05], including only what is directly relevant to the present work. For a more general discussion, see [GGKT05, RTV04].)

1.1. Black-Box Lower Bounds

At the most general level, a construction of a commitment scheme based on one-way permutations may be viewed as a procedure P which takes as input (a description of) a permutation π and outputs (a description of) two circuits $(\mathcal{S}, \mathcal{R})$ (here, \mathcal{S} represents the *sender* while \mathcal{R} represents the *receiver*; see Sections 1.2 and 2.2) realizing the desired commitment functionality whenever π is a permutation. If the construction is *black-box*, this means that P relies only

on the input/output behavior of π and *not* on any internal structure of the implementation of π ; formally, this means that the construction can be described as a pair of oracle procedures $(\mathcal{S}^{\cdot}, \mathcal{R}^{\cdot})$ such that $(\mathcal{S}^{\pi}, \mathcal{R}^{\pi})$ realizes the desired functionality of a commitment scheme for any permutation π .

Besides achieving some desired functionality, a construction of a commitment scheme should also be “secure” in some sense. There are various ways this can be formalized (see [RTV04]); we will be interested here in *weak* black-box constructions which offer the following guarantee:

If π is a one-way permutation, then the scheme $(\mathcal{S}^{\pi}, \mathcal{R}^{\pi})$ is “secure” against all efficient adversaries (who are *not* given oracle access to π),

where “secure” in the above refers to some notions of hiding and binding that will be formalized later. The distinction between whether an adversary is given oracle access to π or not is important since the above should hold *even when π is not efficiently computable* (and so the only way for an efficient adversary to evaluate π , in general, may be via oracle access to π). Note, however, that a weak black-box construction suffices to give implementations with meaningful security guarantees in the real world: in this case, π *will* be efficiently-computable and furthermore an explicit circuit for π will be known; hence, it is irrelevant whether an adversary is given oracle access to π or not. Note also that weak black-box constructions are the *weakest* type of black-box construction considered in [RTV04], and hence impossibility results for weak black-box constructions rule out other black-box constructions as well.

Although most currently-known constructions are black-box, it is important to recognize that a number of non-black-box constructions do exist. As an example, all known constructions of public-key encryption schemes secure against chosen-ciphertext attacks based on trapdoor permutations (e.g., [DDN00]) are non black-box. (See [GGKT05] for additional examples.) Nevertheless, a black-box impossibility result is useful in that it indicates the techniques necessary to achieve a particular result. Furthermore, known non-black-box constructions are much less efficient than black-box ones, and so a black-box impossibility result can be said to rule out “practical” constructions.

1.2. Our Results

With the above in mind, we may now describe our results a bit more formally. An interactive commitment scheme for m -bit messages is a pair of procedures $(\mathcal{S}, \mathcal{R})$ which operates in two phases. In the *commitment phase*, the *sender* \mathcal{S} takes as input a message $M \in \{0, 1\}^m$ and interacts with the *receiver* \mathcal{R} ; we will refer to the view of \mathcal{R} at the conclusion of this phase as the *commitment* to M . In the *decommitment phase*, the sender forwards a *decommitment* to \mathcal{R} which, in particular, reveals M . Without loss of generality, we will assume that the decommitment simply consists of M along with the random coins used by \mathcal{S} during the commitment phase.

A commitment scheme should guarantee both *hiding* and *binding*, where informally these mean that (1) the receiver should have no information about

M before the decommitment phase while (2) the sender should be committed to a unique message at the end of the commitment phase. More formally, a commitment scheme is *statistically binding* if it satisfies the following:

Hiding: For any $M, M' \in \{0, 1\}^m$, the distribution over commitments (by the honest sender \mathcal{S}) to M is computationally indistinguishable from the distribution over commitments (by \mathcal{S}) to M' , even when \mathcal{S} interacts with a malicious (but computationally bounded) receiver \mathcal{R}^* .

(Statistical) binding: The probability (over coin tosses r of the honest receiver \mathcal{R}) that there exist distinct M, M' and coins s, s' for \mathcal{S} such that the corresponding commitments to M, M' are identical is at most ε_b . When $\varepsilon_b = 0$ we say the scheme is *perfectly binding*.

Note that the formulation of the binding requirement ensures security even against an all-powerful sender. Our definition of the binding requirement is somewhat stronger than the usual one which, roughly speaking, requires only that a computationally-unbounded sender *without* knowledge of r be unable to *find* distinct M, M' and coins s, s' such that the corresponding commitments are identical (except with some probability ε_b). For two-round public-coin schemes (where the receiver simply sends a random string and the sender responds with a commitment) and perfectly-binding schemes, however, the notions are identical. Looking ahead, we remark that all the constructions we show in Section 4 satisfy the strong definition of binding given above.

Say a permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is *one-way with security S* if any circuit of size at most S inverts π on at most a fraction $1/S$ of its inputs. Our main result may be stated as follows: any weak black-box construction of a statistically-binding commitment scheme based on a one-way permutation with security S requires $\Omega((m - \log(1 + 2^m \cdot \varepsilon_b)) / \log S)$ invocations of the permutation (by the sender and receiver combined for statistically-binding schemes, and by the sender alone for perfectly-binding schemes). Formally, we show that any construction beating this bound would imply the unconditional existence of a statistically-binding commitment scheme; or, put another way, the only way to develop a more efficient construction of a commitment scheme based on one-way permutations is to construct a commitment scheme *from scratch*. The existence of a commitment scheme implies the existence of one-way functions, and hence $\mathcal{P} \neq \mathcal{NP}$, and so any black-box construction beating our bound would also imply a *proof* that $\mathcal{P} \neq \mathcal{NP}$.

For perfectly-binding schemes, our bound shows that $\Omega(m / \log S)$ invocations of the one-way permutation are needed; our bound in this case matches the efficiency achieved by the construction of Blum [Blu83], instantiated using the Goldreich-Levin hard-core bits of a one-way permutation [GL89]. This is discussed further in Section 4, where we also compare our bounds to known constructions of statistically-binding schemes.

A natural adaptation of our bounds applies also to constructions of commitment schemes based on oracle access to *trapdoor* one-way permutations (see [GGKT05] for definitions).

2. Definitions

2.1. Preliminaries

Let A^f denote a circuit A with oracle access to the function f . A function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is (S, ε) -one-way if for every circuit A of size at most S we have

$$\Pr_x[A^f(f(x)) \in f^{-1}(f(x))] \leq \varepsilon.$$

To reduce the number of parameters, we will call a function S -hard if it is $(S, 1/S)$ -one way.

Let Π_t denote the set of all permutations over $\{0, 1\}^t$. We rely on the following result:

Theorem 1 ([GGKT05]). *For sufficiently large t , a random $\pi \in \Pi_t$ is $2^{t/5}$ -hard with probability at least $1 - 2^{-2^{t/2}}$.*

Let $a||b$ denote the concatenation of strings a and b . For $t < n$, let $\Pi_{t,n}$ denote the subset of Π_n such that $\pi \in \Pi_{t,n}$ iff $\pi(a||b) = \hat{\pi}(a)||b$ for some $\hat{\pi} \in \Pi_t$ (i.e., the last $n - t$ bits of the input are fixed). A corollary of Theorem 1 is that if $t = 5 \log S$, then for any $n > t$ a randomly chosen $\pi \in \Pi_{t,n}$ is S -hard with high probability; more formally:

Corollary 2 ([GGKT05]). *For sufficiently large t and $n > t$, a random $\pi \in \Pi_{t,n}$ is $2^{t/5}$ -hard with probability at least $1 - 2^{-2^{t/2}}$.*

We say that two distributions \mathcal{X}, \mathcal{Y} are (S, ε) -indistinguishable, and write $\mathcal{X} \stackrel{(S, \varepsilon)}{\approx} \mathcal{Y}$, if for every circuit Dist of size at most S , we have

$$\left| \Pr_{x \in \mathcal{X}}[\text{Dist}(x) = 1] - \Pr_{x \in \mathcal{Y}}[\text{Dist}(x) = 1] \right| \leq \varepsilon.$$

2.2. Commitment Schemes

A *commitment scheme* for m -bit messages is defined by a pair of probabilistic, interactive algorithms $(\mathcal{S}, \mathcal{R})$ representing a *sender* and a *receiver*, respectively. (We remark that $(\mathcal{S}, \mathcal{R})$ describe the commitment phase only; recall that, without loss of generality, we will assume that the sender simply reveals M and its random tape s in order to decommit.) The inputs to \mathcal{S} are a message $M \in \{0, 1\}^m$ and a random tape s , while the input to \mathcal{R} is a random tape r . Let $\langle \mathcal{S}(M; s), \mathcal{R}^*(r) \rangle$ denote the view of a (possibly malicious) receiver \mathcal{R}^* following an interaction with the sender on the specified inputs; this view consists of the receiver's randomness and the messages it receives from the sender during the interaction. (When the receiver makes queries to an oracle, the view also includes the answers it receives from this oracle.) For a message M and receiver \mathcal{R}^* , define

$$\langle \mathcal{S}(M), \mathcal{R}^* \rangle \stackrel{\text{def}}{=} \{s, r \leftarrow \{0, 1\}^* : \langle \mathcal{S}(M; s), \mathcal{R}^*(r) \rangle\};$$

i.e., this denotes the distribution over the view of \mathcal{R}^* following an interaction with the honest sender who is committing to message M .

We now define the security of a commitment scheme. In this paper we only deal with statistically binding schemes, as reflected in the definitions that follow.

Definition 3. Let $(\mathcal{S}, \mathcal{R})$ be a commitment scheme for m -bit messages. We say that $(\mathcal{S}, \mathcal{R})$ is (S_h, ε_h) -*hiding* if for every circuit \mathcal{R}^* of size at most S_h and for all $M_0, M_1 \in \{0, 1\}^m$, we have

$$\langle \mathcal{S}(M_0), \mathcal{R}^* \rangle \stackrel{(S_h, \varepsilon_h)}{\approx} \langle \mathcal{S}(M_1), \mathcal{R}^* \rangle. \quad (1)$$

(To be meaningful, S_h should be at least the size of the honest receiver algorithm \mathcal{R} .) We say that $(\mathcal{S}, \mathcal{R})$ is ε_b -*binding* if

$$\Pr_r \left[\exists \text{ distinct } M, M' \in \{0, 1\}^m, s, s' \text{ such that } \langle \mathcal{S}(M'; s'), \mathcal{R}(r) \rangle = \langle \mathcal{S}(M; s), \mathcal{R}(r) \rangle \right] \leq \varepsilon_b.$$

Note that if a commitment scheme is ε_b -binding then even an all-powerful sender cannot commit in such a way that it can later decommit to two different messages, except with probability (at most) ε_b . We say that $(\mathcal{S}, \mathcal{R})$ is ε_b -*binding for an honest sender* if for all $M \in \{0, 1\}^m$, we have

$$\Pr_{s,r} \left[\exists M' \in \{0, 1\}^m \setminus M, s' \text{ such that } \langle \mathcal{S}(M'; s'), \mathcal{R}(r) \rangle = \langle \mathcal{S}(M; s), \mathcal{R}(r) \rangle \right] \leq \varepsilon_b.$$

Roughly speaking, such a scheme satisfies the following property: if the sender is honest during the commitment phase (and uses a pre-fixed message M and truly random coins s) then the sender cannot later decommit to a different value M' except with probability (at most) ε_b . If $\varepsilon_b = 0$ in either of the above definitions, we say the scheme is *perfectly binding* (resp., *perfectly binding for an honest sender*).

$(\mathcal{S}, \mathcal{R})$ is $(S_h, \varepsilon_h, \varepsilon_b)$ -*secure* (resp., *secure for an honest sender*) if $(\mathcal{S}, \mathcal{R})$ is (S_h, ε_h) -hiding and ε_b -binding (resp., binding for an honest sender).

We may now define a weak black-box construction of a commitment scheme based on one-way permutations.

Definition 4. A construction of a commitment scheme for m -bit messages based on one-way permutations is a pair of oracle algorithms $(\mathcal{S}^{(\cdot)}, \mathcal{R}^{(\cdot)})$ such that, for all $\pi \in \Pi_n$, the resulting $(\mathcal{S}^\pi, \mathcal{R}^\pi)$ is a commitment scheme for m -bit messages. We say that $(\mathcal{S}^{(\cdot)}, \mathcal{R}^{(\cdot)})$ is $(S_p, S_h, \varepsilon_h, \varepsilon_b)$ -*secure* (resp., *secure for an honest sender*) if $(\mathcal{S}^\pi, \mathcal{R}^\pi)$ is ε_b -binding (resp., binding for an honest sender) for every $\pi \in \Pi_n$, and furthermore for every $\pi \in \Pi_n$ that is S_p -hard, scheme $(\mathcal{S}^\pi, \mathcal{R}^\pi)$ is (S_h, ε_h) -hiding.³

³Our constructions are *weak* black-box in the sense that the distinguisher (implicit in (1)) is not given oracle access to π .

2.3. Pairwise-Independent Function Families

Let H be a family of functions mapping m -bit strings to m' -bit strings. We assume that the following can be done in time polynomial in m : (1) selecting a function $h \in H$ uniformly at random; and (2) given $h \in H$ and $x \in \{0, 1\}^m$, evaluating $h(x)$. We say H is a *pairwise-independent hash family* [CW79] if for any distinct $x_1, x_2 \in \{0, 1\}^m$ and any $y_1, y_2 \in \{0, 1\}^{m'}$ we have:

$$\Pr_{h \in H} [h(x_1) = y_1 \wedge h(x_2) = y_2] = 2^{-2m'}.$$

Constructions satisfying the above requirements are well known.

3. Lower Bounding the Efficiency of Commitment

Let $(\mathcal{S}^{(\cdot)}, \mathcal{R}^{(\cdot)})$ be an $(S_p, S_h, \varepsilon_h, \varepsilon_b)$ -secure construction of a commitment scheme for m -bit messages (based on one-way permutations). For $\varepsilon_b > 0$, we prove that unless \mathcal{S} and \mathcal{R} (combined) make $\Omega((m - \log(1 + 2^m \cdot \varepsilon_b)) / \log S_p)$ queries to their oracle, there exists (constructively) a commitment scheme $(\tilde{\mathcal{S}}, \tilde{\mathcal{R}})$ secure for an honest sender that does not require any oracle access *at all* (i.e., the scheme is secure unconditionally). For $\varepsilon_b = 0$, we show a similar result but where the implication holds unless \mathcal{S} alone makes $\Omega(m / \log S_p)$ queries to its oracle. In either case, by applying a result of Impagliazzo and Luby [IL89] (cf. also Lemma 5 below) this implies the unconditional existence of a one-way function, which in turn can be used to give an unconditional construction of a commitment scheme [Naor91].

We describe here some of the intuition behind our proof, focusing for ease of exposition on the case that $(\mathcal{S}^{(\cdot)}, \mathcal{R}^{(\cdot)})$ is perfectly binding. As in [GGKT05], our starting point is that a random $\pi \in \Pi_{t,n}$ (for $t = \Theta(\log S_p)$) is S_p -hard with all but negligible probability (cf. Corollary 2). Consider the non-interactive scheme $(\mathcal{S}', \mathcal{R}')$ in which \mathcal{S}' locally runs $(\mathcal{S}^{(\cdot)}, \mathcal{R}^{(\cdot)})$, simulating a random $\pi \in \Pi_{t,n}$ for the algorithms at hand,⁴ and then sends the resulting view of \mathcal{R} to \mathcal{R}' . Decommithment, as usual, is performed by having \mathcal{S}' send the message and all the random coins it used to \mathcal{R}' .

It is quite straightforward to show that $(\mathcal{S}', \mathcal{R}')$ still satisfies hiding. Binding, however, may not necessarily hold even when \mathcal{S}' is honest during the commitment phase. To see the issue, assume \mathcal{S}' commits to a message M using coins s for $\mathcal{S}^{(\cdot)}$, coins r for $\mathcal{R}^{(\cdot)}$, and coins y to simulate the permutation. Let C denote the resulting view of \mathcal{R} , and let P denote the set of t -bit query/answer prefixes made by \mathcal{S} during the computation. To claim binding, we would need to argue that there does not exist a message $M' \neq M$ along with coins s', y' , with an associated set of query/answer prefixes P' , that produce an identical view C (note that the coins r are fixed, since r is explicit in the view C that was already sent to \mathcal{R}'). The most we can claim, though, is that this is true *as long*

⁴This can be done easily by selecting random t -bit answer-prefixes for any new t -bit query-prefixes, as needed; see details in the proof of Theorem 6.

as $P' = P$, since binding is only guaranteed to hold when the permutation π is fixed, but not when the sender can “change” the permutation after the fact.

What we *can* show is that a weaker form of (honest-sender) binding holds for $(\mathcal{S}', \mathcal{R}')$. Observe that for any possible P' (as defined above), there is at most *one* message M' to which the sender can successfully decommit by sending M', s', y' with associated query/answer set P' ; this is because $(\mathcal{S}^{(\cdot)}, \mathcal{R}^{(\cdot)})$ is perfectly binding for any *fixed* permutation. But this implies that there are at most $2^{2t|P'|} = 2^{2tq}$ different messages to which the sender can successfully decommit, where q is the total number of queries made by \mathcal{S} (note that the oracle queries/answers of \mathcal{R} are already fixed by the view C). Although this clearly violates binding, it does somewhat limit the space of possible messages to which the sender can decommit as long as $2^{2tq} < 2^m$.

We next show how to “bootstrap” from the weak form of binding achieved by $(\mathcal{S}', \mathcal{R}')$ to construct a non-interactive scheme $(\bar{\mathcal{S}}, \bar{\mathcal{R}})$ that achieves “full” binding (for an honest sender) with noticeable probability. Sender $\bar{\mathcal{S}}$, on input a message M , proceeds as follows: it first chooses a function h uniformly at random from a pairwise-independent hash family mapping m -bit strings to m -bit strings. It then computes the views $C_1 = \mathcal{S}'(M)$, $C_2 = \mathcal{S}'(h(M))$, and sends $C_1 \| C_2 \| h$ to $\bar{\mathcal{R}}$. Hiding for this scheme follows easily via a standard hybrid argument and relying on the fact that $(\mathcal{S}', \mathcal{R}')$ is hiding. As for binding (for an honest sender), we have already seen that C_1 can be decommitted to a set S_1 of at most $2^{2tq} < 2^m$ different messages, and similarly C_2 can be decommitted to a set S_2 of at most 2^{2tq} different messages. For binding not to hold, there must exist an $M' \neq M$ with $M' \in S_1$ and $h(M') \in S_2$. Using the pairwise-independence of h , we can argue that this occurs with only “small” probability over choice of h . Thus, binding for $(\bar{\mathcal{S}}, \bar{\mathcal{R}})$ (for an honest sender) holds with noticeable probability.

3.1. A Technical Lemma

We begin by showing that the existence of a commitment scheme secure for honest senders implies the existence of a one-way function. Although the result can be derived from [IL89], we give a simple and more direct proof here.

Lemma 5. *Let $(\mathcal{S}, \mathcal{R})$ be a commitment scheme for m -bit messages which is $(S_h, \varepsilon_h, \varepsilon_b)$ -secure for an honest sender. Let $S_{\mathcal{S}}, S_{\mathcal{R}}$ be the sizes of the circuits computing \mathcal{S}, \mathcal{R} , respectively. Then there exists an $(S_h - S_{\mathcal{S}} + S_{\mathcal{R}} - O(m), \varepsilon_h + 2\varepsilon_b)$ -one-way function.*

Proof. Let $S^* = S_h - S_{\mathcal{S}} + S_{\mathcal{R}} - O(m)$ and $\varepsilon^* = \varepsilon_h + 2\varepsilon_b$. Define a function f via $f(M, s, r) \stackrel{\text{def}}{=} \langle \mathcal{S}(M; s), \mathcal{R}(r) \rangle$. We claim that f is (S^*, ε^*) -one-way. Assume the contrary. Then there exists a circuit B of size at most S^* such that

$$\text{Succ}_{B,f}^{\text{owf}} \stackrel{\text{def}}{=} \Pr_{M,s,r} [B(f(M, s, r)) \in f^{-1}(f(M, s, r))] > \varepsilon^*.$$

We use B to construct a circuit A that violates the hiding property of $(\mathcal{S}, \mathcal{R})$. On input (M_0, M_1, C) , where C is either a commitment to M_0 or M_1 , A computes

$(M', s', r') \leftarrow B(C)$ and checks whether $f(M', s', r') \stackrel{?}{=} C$ and whether $M' \stackrel{?}{=} M_0$. If both hold, A outputs 0; otherwise, it outputs 1. Note that $|A| = |B| + S_S + S_{\mathcal{R}} + O(m) \leq S_h$.

Let $\text{Bad} \stackrel{\text{def}}{=} \{(M, s, r) \mid \exists M' \neq M, s' : \langle \mathcal{S}(M; s), \mathcal{R}(r) \rangle = \langle \mathcal{S}(M'; s'), \mathcal{R}(r) \rangle\}$. In what follows, note that if $(M', s', r') \in f^{-1}(f(M, s, r))$ then $r' = r$, as r is included in the receiver's view. We have:

$$\begin{aligned}
& \Pr_{\substack{M_0, M_1 \\ C \in \langle \mathcal{S}(M_0), \mathcal{R} \rangle}} [A(M_0, M_1, C) = 0] \\
&= \Pr_{M_0, s, r} \left[\begin{array}{l} (M', s', r') \leftarrow B(f(M_0, s, r)) : \\ (M', s', r') \in f^{-1}(f(M_0, s, r)) \wedge M' = M_0 \end{array} \right] \\
&\geq \Pr_{M_0, s, r} \left[\begin{array}{l} (M', s', r') \leftarrow B(f(M_0, s, r)) : \\ (M', s', r') \in f^{-1}(f(M_0, s, r)) \wedge (M_0, s, r) \notin \text{Bad} \end{array} \right] \\
&\geq \Pr_{M_0, s, r} \left[\begin{array}{l} (M', s', r') \leftarrow B(f(M_0, s, r)) : \\ (M', s', r') \in f^{-1}(f(M_0, s, r)) \end{array} \right] - \Pr_{M_0, s, r} [(M_0, s, r) \in \text{Bad}] \\
&\geq \text{Succ}_{B, f}^{\text{owf}} - \varepsilon_b \\
&= \varepsilon_h + \varepsilon_b.
\end{aligned}$$

Furthermore, we have:

$$\begin{aligned}
& \Pr_{\substack{M_0, M_1 \\ C \in \langle \mathcal{S}(M_1), \mathcal{R} \rangle}} [A(M_0, M_1, C) = 0] \\
&= \Pr_{\substack{M_0, M_1 \\ s, r}} \left[\begin{array}{l} (M', s', r') \leftarrow B(f(M_1, s, r)) : \\ (M', s', r') \in f^{-1}(f(M_1, s, r)) \wedge M' = M_0 \end{array} \right] \\
&\leq \Pr_{\substack{M_0, M_1 \\ s, r}} \left[\begin{array}{l} (M', s', r') \leftarrow B(f(M_1, s, r)) : \\ (M', s', r') \in f^{-1}(f(M_1, s, r)) \wedge (M_1, s, r) \in \text{Bad} \end{array} \right] \\
&\leq \Pr_{M_1, s, r} [(M_1, s, r) \in \text{Bad}] \\
&\leq \varepsilon_b.
\end{aligned}$$

Putting everything together, we have:

$$\left| \Pr_{\substack{M_0, M_1 \\ C \in \langle \mathcal{S}(M_0), \mathcal{R} \rangle}} [A(M_0, M_1, C) = 0] - \Pr_{\substack{M_0, M_1 \\ C \in \langle \mathcal{S}(M_1), \mathcal{R} \rangle}} [A(M_0, M_1, C) = 0] \right| > \varepsilon_h.$$

But this implies that there exist two messages M_0, M_1 for which A can distinguish $\langle \mathcal{S}(M_0), \mathcal{R} \rangle$ from $\langle \mathcal{S}(M_1), \mathcal{R} \rangle$ with probability greater than ε_h , contradicting the hiding of $(\mathcal{S}, \mathcal{R})$. \square

3.2. Main Result

We now formalize the intuition that was discussed earlier. The proof below is not as straightforward as the intuition would suggest, since some technical work is required to deal with the case of statistical (as opposed to perfect) binding.

Theorem 6. Let $(\mathcal{S}^{(\cdot)}, \mathcal{R}^{(\cdot)})$ be an $(S_p, S_h, \varepsilon_h, \varepsilon_b)$ -secure construction of a commitment scheme for m -bit messages that expects an oracle $\pi \in \Pi_n$. Let $t = 5 \log S_p$. Assume $\varepsilon_h \leq 1/8 - 2^{1-S_p}$. If $\varepsilon_b > 0$ and \mathcal{S} and \mathcal{R} make a total of $q \leq (m-2 - \log(1+2^{m+1} \cdot \varepsilon_b))/4t$ queries to their oracle, or if $\varepsilon_b = 0$ and \mathcal{S} makes $q_S \leq (m-2)/4t$ queries to its oracle, then there exists a commitment scheme (without access to any oracle) for m -bit messages which is $(S_h, 1/4, 1/4)$ -secure for an honest sender.

Applying Lemma 5, the conclusion of the theorem implies the existence of a one-way function (without access to any oracle).

Proof. We construct a commitment scheme $(\bar{\mathcal{S}}, \bar{\mathcal{R}})$ for m -bit messages, following the intuition outlined earlier. The construction makes use of a procedure $\mathcal{S}\mathcal{I}\mathcal{M}$ that simulates a random permutation in $\Pi_{t,n}$ as follows: $\mathcal{S}\mathcal{I}\mathcal{M}$ maintains a list L which is initially empty. To respond to a query $a||a'$, where $|a| = t$ and $|a'| = n - t$, procedure $\mathcal{S}\mathcal{I}\mathcal{M}$ first checks whether there exists a value b such that $(a, b) \in L$. If so, $\mathcal{S}\mathcal{I}\mathcal{M}$ returns $b||a'$. Otherwise, it picks $b \in \{0, 1\}^t \setminus \{\hat{b} \mid \exists \hat{a} : (\hat{a}, \hat{b}) \in L\}$ uniformly at random, adds (a, b) to L , and returns $b||a'$. We let $\mathcal{S}\mathcal{I}\mathcal{M}_y$ denote an execution of $\mathcal{S}\mathcal{I}\mathcal{M}$ using random coins y .

Let H be a pairwise-independent family of functions from m -bit strings to m -bit strings. Define $\bar{\mathcal{S}}$ as follows. On input a message $M \in \{0, 1\}^m$, $\bar{\mathcal{S}}$ chooses uniformly at random $h \in H$ and values $s_1, r_1, y_1, s_2, r_2, y_2$. It then computes⁵ $C_1 = \langle \mathcal{S}^{\mathcal{S}\mathcal{I}\mathcal{M}_{y_1}}(M; s_1), \mathcal{R}^{\mathcal{S}\mathcal{I}\mathcal{M}_{y_1}}(r_1) \rangle$ and $C_2 = \langle \mathcal{S}^{\mathcal{S}\mathcal{I}\mathcal{M}_{y_2}}(h(M); s_2), \mathcal{R}^{\mathcal{S}\mathcal{I}\mathcal{M}_{y_2}}(r_2) \rangle$, and outputs $C_1||C_2||h$. Decommitment, as usual, is done by having $\bar{\mathcal{S}}$ reveal M and all the random coins used during the commitment phase. The claim that $(\bar{\mathcal{S}}, \bar{\mathcal{R}})$ is $(S_h, 1/4, 1/4)$ -secure for an honest sender follows from the next two lemmas.

Lemma 7. $(\bar{\mathcal{S}}, \bar{\mathcal{R}})$ is $(S_h, 1/4)$ -hiding.

Proof. The proof is quite straightforward. The hiding property of $(\mathcal{S}^{(\cdot)}, \mathcal{R}^{(\cdot)})$ guarantees that for any $\pi \in \Pi_n$ that is S_p -hard, for any circuit Dist of size at most S_h , and for any messages $M_0, M_1 \in \{0, 1\}^m$, we have

$$\left| \Pr_{C \in \langle \mathcal{S}^\pi(M_0), \mathcal{R}^\pi \rangle} [\text{Dist}(C) = 0] - \Pr_{C \in \langle \mathcal{S}^\pi(M_1), \mathcal{R}^\pi \rangle} [\text{Dist}(C) = 0] \right| \leq \varepsilon_h.$$

To save on notation, let $\text{Com}^\pi(M)$ denote $\langle \mathcal{S}^\pi(M_0), \mathcal{R}^\pi \rangle$; i.e., $\text{Com}^\pi(M)$ denotes the distribution on views of the honest receiver when \mathcal{S} commits to M and the parties are using oracle π . A straightforward hybrid argument shows that for any $\pi_1, \pi_2 \in \Pi_n$ that are S_p -hard, for any circuit Dist of size at most S_h , and

⁵The permutations simulated by $\mathcal{S}\mathcal{I}\mathcal{M}$ in the computations of C_1, C_2 will, in general, be different. The theorem can be strengthened (improving the bound on ε_h) by having $\mathcal{S}\mathcal{I}\mathcal{M}$ provide a consistent simulation for both computations. We forgo this for simplicity.

for any $M_0, M_1 \in \{0, 1\}^m$, we have

$$2\varepsilon_h \geq \left| \Pr_{\substack{h \in H \\ C_1 \in \text{Com}^{\pi_1}(M_0) \\ C_2 \in \text{Com}^{\pi_2}(h(M_0))}} [\text{Dist}(C_1 \| C_2 \| h) = 0] - \Pr_{\substack{h \in H \\ C_1 \in \text{Com}^{\pi_1}(M_1) \\ C_2 \in \text{Com}^{\pi_2}(h(M_1))}} [\text{Dist}(C_1 \| C_2 \| h) = 0] \right|.$$

Corollary 2 shows that a random permutation $\pi \in \Pi_{t,n}$ is S_p -hard except with probability at most $2^{-S_p^{5/2}} \leq 2^{-S_p}$. Using a union bound and a simple averaging argument, we see that for any circuit Dist of size at most S_h and for any $M_0, M_1 \in \{0, 1\}^m$,

$$2\varepsilon_h + 2^{1-S_p} \geq \left| \Pr_{\substack{\pi_1, \pi_2 \in \Pi_{t,n} \\ h \in H \\ C_1 \in \text{Com}^{\pi_1}(M_0) \\ C_2 \in \text{Com}^{\pi_2}(h(M_0))}} [\text{Dist}(C_1 \| C_2 \| h) = 0] - \Pr_{\substack{\pi_1, \pi_2 \in \Pi_{t,n} \\ h \in H \\ C_1 \in \text{Com}^{\pi_1}(M_1) \\ C_2 \in \text{Com}^{\pi_2}(h(M_1))}} [\text{Dist}(C_1 \| C_2 \| h) = 0] \right|.$$

Since \mathcal{SIM} perfectly simulates a random $\pi \in \Pi_{t,n}$, this is exactly equivalent to

$$\left| \Pr_{C \in \langle \bar{\mathcal{S}}(M_0), \mathcal{R}^* \rangle} [\text{Dist}(C) = 0] - \Pr_{C \in \langle \bar{\mathcal{S}}(M_1), \mathcal{R}^* \rangle} [\text{Dist}(C) = 0] \right| \leq 2\varepsilon_h + 2^{1-S_p} \leq 1/4$$

for any \mathcal{R}^* and any circuit Dist of size at most S_h , where the last inequality uses the assumption that $\varepsilon_h \leq 1/8 - 2^{1-S_p}$. The hiding property therefore holds as claimed. \square

Lemma 8. $(\bar{\mathcal{S}}, \bar{\mathcal{R}})$ is 1/4-binding for an honest sender.

Proof. For ease of notation, let

$$\text{Com}(M, s, r, y) \stackrel{\text{def}}{=} \langle \mathcal{S}^{\text{SIM}_y}(M; s), \mathcal{R}^{\text{SIM}_y}(r) \rangle.$$

Fix an arbitrary $M \in \{0, 1\}^m$. We are interested in the following probability:

$$\begin{aligned} \text{NoBind} &\stackrel{\text{def}}{=} \Pr_{\bar{s}} \left[\exists M' \in \{0, 1\}^m \setminus M, \bar{s}' \text{ such that } \langle \bar{\mathcal{S}}(M'; \bar{s}'), \bar{\mathcal{R}} \rangle = \langle \bar{\mathcal{S}}(M; \bar{s}), \bar{\mathcal{R}} \rangle \right] \\ &= \Pr_{\substack{h \in H \\ s_1, r_1, y_1 \\ s_2, r_2, y_2}} \left[\begin{array}{l} \exists M' \in \{0, 1\}^m \setminus M, h', s'_1, r'_1, y'_1, s'_2, r'_2, y'_2 \text{ such that} \\ \text{Com}(M', s'_1, r'_1, y'_1) \| \text{Com}(h'(M'), s'_2, r'_2, y'_2) \| h' \\ = \text{Com}(M, s_1, r_1, y_1) \| \text{Com}(h(M), s_2, r_2, y_2) \| h \end{array} \right] \\ &= \Pr_{\substack{h \in H \\ s_1, r_1, y_1 \\ s_2, r_2, y_2}} \left[\begin{array}{l} \exists M' \in \{0, 1\}^m \setminus M, s'_1, y'_1, s'_2, y'_2 \text{ such that} \\ \text{Com}(M', s'_1, r_1, y'_1) = \text{Com}(M, s_1, r_1, y_1) \wedge \\ \text{Com}(h(M'), s'_2, r_2, y'_2) = \text{Com}(h(M), s_2, r_2, y_2) \end{array} \right], \end{aligned}$$

where in the last equality we use the fact that h', r'_1, r'_2 and h, r_1, r_2 are explicit in the view of \mathcal{R} . Letting

$$\text{Decom}(M, s, r, y) \stackrel{\text{def}}{=} \left\{ M' \in \{0, 1\}^m \mid \begin{array}{c} \exists s', y' \text{ such that} \\ \text{Com}(M', s', r, y') = \text{Com}(M, s, r, y) \end{array} \right\},$$

we may write:

$$\text{NoBind} = \Pr_{\substack{h \in H \\ s_1, r_1, y_1 \\ s_2, r_2, y_2}} \left[\begin{array}{c} \exists M' \in \{0, 1\}^m \setminus M \text{ such that} \\ M' \in \text{Decom}(M, s_1, r_1, y_1) \wedge \\ h(M') \in \text{Decom}(h(M), s_2, r_2, y_2) \end{array} \right].$$

For any integer q , let Perm_t^q denote the set of “partial permutations” of size q over t -bit strings; formally, Perm_t^q contains all sets $P \subseteq \{0, 1\}^t \times \{0, 1\}^t$ such that P contains exactly q tuples and such that for all a there exists at most one b with $(a, b) \in P$ and for all b there exists at most one a such that $(a, b) \in P$. (I.e., P can be extended to a permutation over $\{0, 1\}^t$.) Let $q_{\mathcal{S}}$ (resp., $q_{\mathcal{R}}$) denote the number of queries made by \mathcal{S} (resp., \mathcal{R}) to its oracle,⁶ and let $q = q_{\mathcal{S}} + q_{\mathcal{R}}$. Let $\text{queries}(M, s, r, y) \in \text{Perm}_t^q$ denote the set of query/answer prefixes made by either \mathcal{S} or \mathcal{R} to SIM during the computation of $\text{Com}(M, s, r, y)$ (i.e., $(a, b) \in \text{queries}(M, s, r, y)$ iff an oracle query $a \| a'$, by either \mathcal{S} or \mathcal{R} , is answered by SIM with $b \| a'$ during the computation of $\text{Com}(M, s, r, y)$). Define $\text{queries}_{\mathcal{S}}(M, s, r, y)$ (resp., $\text{queries}_{\mathcal{R}}(M, s, r, y)$) similarly, where this refers exclusively to queries made by \mathcal{S} (resp., \mathcal{R}).

Define r as *good for* $P \in \text{Perm}_t^q$ if there do not exist distinct M', M'' , along with s', s'', y', y'' , such that

- $\text{Com}(M', s', r, y') = \text{Com}(M'', s'', r, y'')$; and
- $\text{queries}(M', s', r, y') = \text{queries}(M'', s'', r, y'') = P$.

Say r is *good* if it is good for all $P \in \text{Perm}_t^q$.

We first observe that for a good r , the set $\text{Decom}(M, s, r, y)$ contains at most $|\text{Perm}_t^{q_{\mathcal{S}}}| < 2^{2tq_{\mathcal{S}}}$ messages. Otherwise, by the pigeonhole principle, there exists a $P_{\mathcal{S}} \in \text{Perm}_t^{q_{\mathcal{S}}}$ and distinct messages $M', M'' \in \text{Decom}(M, s, r, y)$, along with s', s'', y', y'' , such that $\text{Com}(M', s', r, y') = \text{Com}(M, s, r, y) = \text{Com}(M'', s'', r, y'')$ and $\text{queries}_{\mathcal{S}}(M', s', r, y') = \text{queries}_{\mathcal{S}}(M'', s'', r, y'') = P_{\mathcal{S}}$. (Notice also that $\text{queries}_{\mathcal{R}}(M', s', r, y') = \text{queries}_{\mathcal{R}}(M, s, r, y) = \text{queries}_{\mathcal{R}}(M'', s'', r, y'')$, as these queries are explicit in the receiver’s views.) But then r is not good for $P \stackrel{\text{def}}{=} P_{\mathcal{S}} \cup \text{queries}_{\mathcal{R}}(M, s, r, y)$, contradicting the assumption that r is good.

Fix some $P \in \text{Perm}_t^q$, and let π_P denote an arbitrary extension of P to a

⁶Without loss of generality, we assume exactly $q_{\mathcal{S}}$ (resp., $q_{\mathcal{R}}$) queries are always made.

permutation in $\Pi_{t,n}$. We have

$$\begin{aligned}
& \Pr_r[r \text{ is not good for } P] \\
&= \Pr_r \left[\begin{array}{l} \exists \text{ distinct } M', M'' \text{ and } s', s'', y', y'' \text{ such that} \\ \text{Com}(M', s', r, y') = \text{Com}(M'', s'', r, y'') \wedge \\ \text{queries}(M', s', r, y') = \text{queries}(M'', s'', r, y'') = P \end{array} \right] \\
&\leq \Pr_r \left[\begin{array}{l} \exists \text{ distinct } M', M'' \text{ and } s', s'' \text{ such that} \\ \langle \mathcal{S}^{\pi_P}(M'; s'), \mathcal{R}^{\pi_P}(r) \rangle = \langle \mathcal{S}^{\pi_P}(M''; s''), \mathcal{R}^{\pi_P}(r) \rangle \end{array} \right] \\
&\leq \varepsilon_b,
\end{aligned}$$

by the binding property of $(\mathcal{S}^{(\cdot)}, \mathcal{R}^{(\cdot)})$. Applying a union bound over all elements of Perm_t^q , we obtain:

$$\Pr_r[r \text{ is not good}] < 2^{2tq} \cdot \varepsilon_b.$$

We proceed to bound **NoBind**. We have:

$$\text{NoBind} \leq \underbrace{\Pr_{\substack{h \in H \\ s_1, r_1, y_1 \\ s_2, r_2, y_2}} \left[\begin{array}{l} \exists M' \in \{0, 1\}^m \setminus M \text{ such that} \\ M' \in \text{Decom}(M, s_1, r_1, y_1) \\ h(M') \in \text{Decom}(h(M), s_2, r_2, y_2) \end{array} \middle| r_1, r_2 \text{ good} \right]}_{\text{LeftTerm}} + 2^{2tq+1} \cdot \varepsilon_b,$$

where the right term above represents an upper-bound on the probability that either r_1 or r_2 is not good. Continuing with the left term, we have

LeftTerm

$$\begin{aligned}
&= \sum_{M_2 \in \{0, 1\}^m} \Pr_{\substack{h \in H \\ s_1, r_1, y_1 \\ s_2, r_2, y_2}} \left[\begin{array}{l} \exists M' \in \text{Decom}(M, s_1, r_1, y_1) \setminus M \\ \text{and } M'_2 \in \text{Decom}(M_2, s_1, r_1, y_1) \\ \text{such that } h(M) = M_2 \wedge h(M') = M'_2 \end{array} \middle| r_1, r_2 \text{ good} \right] \\
&= \sum_{M_2 \in \{0, 1\}^m} \left(2^{-2m} \cdot \max_{\substack{s_1, y_1 \\ s_2, y_2 \\ \text{good } r_1, r_2}} \left\{ |\text{Decom}(M, s_1, r_1, y_1)| \cdot |\text{Decom}(M_2, s_2, r_2, y_2)| \right\} \right),
\end{aligned}$$

using pairwise independence of H . Applying the bound on the size of $\text{Decom}(M, s, r, y)$ when r is good, we obtain

$$\text{LeftTerm} \leq 2^{-2m} \cdot 2^m \cdot 2^{4tqs} = 2^{4tqs-m}.$$

Putting everything together, we have

$$\text{NoBind} \leq 2^{4tqs-m} + 2^{2tq+1} \cdot \varepsilon_b.$$

If $\varepsilon_b = 0$ and $q_S \leq (m-2)/4t$, it is easy to see that $\text{NoBind} \leq 1/4$. When $\varepsilon_b > 0$ and $q \leq (m-2-\log(1+2^{m+1} \cdot \varepsilon_b))/4t$, then $2^{4tqs-m} + 2^{2tq+1} \cdot \varepsilon_b \leq 2^{4tq} \cdot (2^{-m} + 2\varepsilon_b)$ and hence $\text{NoBind} \leq 1/4$ in this case as well. The claim follows. \square

This completes the proof of the theorem. \square

4. Upper Bounds on the Efficiency of Commitment

Here, we briefly describe upper bounds on the efficiency of black-box constructions of commitment schemes based on one-way permutations.

4.1. Perfectly-Binding Commitment

A perfectly-binding commitment scheme can be constructed from one-way permutations using the approach of Blum [Blu83] along with the Goldreich-Levin hard-core function paradigm [GL89]. Specifically, let $h : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ be a *hard-core function* (see [Gol01]) for a one-way permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$. To commit to a message $M \in \{0, 1\}^m$, the sender first divides M into $t = \lceil m/\ell \rceil$ blocks N_1, \dots, N_t , each of length ℓ . Then, for each block N_i the sender chooses a random $s_i \in \{0, 1\}^n$ and sends $\pi(s_i), h(s_i) \oplus N_i$ to the receiver. Since there exists a hard-core function with $\ell = O(\log S)$ for any S -hard π (and large enough n) [GL89] (see also [Gol01, Section 2.5.3]), this construction requires $O(m/\log S)$ invocations of π , matching our bound.

4.2. Statistically-Binding Commitment for Single-Bit Messages

Naor [Naor91] showed a construction of a statistically-binding commitment scheme for single-bit messages based on one-way *functions*. Let $G : \{0, 1\}^n \rightarrow \{0, 1\}^{n+k}$ be a pseudorandom generator. The receiver first chooses a random $r \in \{0, 1\}^{n+k}$ and sends this value to the other party. The sender then commits to a bit b as follows: it chooses a random $s \in \{0, 1\}^n$ and sends $G(s)$ if $b = 0$ and $G(s) \oplus r$ if $b = 1$. This scheme is binding with $\varepsilon_b < 2^{2n}/2^{n+k} = 2^{n-k}$.

Although a pseudorandom generator G can be constructed from any one-way function [HILL99], we examine the efficiency of the above scheme when G is based on an S -hard one-way permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ so as to compare the efficiency of the scheme to our bound. In this case, evaluating G requires $O(k/\log S)$ invocations of π [Yao82, BM84, GL89]. Viewing n as fixed, this is $O(\log \varepsilon_b^{-1}/\log S)$ invocations of π (for k polynomial in n).

4.3. Statistically-Binding Constructions for Longer Messages

There are a number of ways to extend the Naor scheme described above for the case of m -bit messages. One obvious approach is to simply run the basic Naor scheme in parallel for each bit of the message, having the sender/receiver use the same value r for all these commitments. This gives a scheme which is binding with $\varepsilon_b < 2^{n-k}$ as before, but where the number of invocations of π required is now $O(mk/\log S)$.

A better approach, suggested in [Naor91], is to have the sender use the above idea to commit to an n -bit *seed* s , and then additionally send $G'(s) \oplus M$ (where M is the sender's message and $G' : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is another pseudorandom generator). This is still binding with $\varepsilon_b < 2^{n-k}$ as before; the number of invocations of π required, however, is $O(nk/\log S + (m-n)/\log S)$ which is more efficient than the previous approach when $m > n$.

A third approach, suggested in [Naor91] as well, utilizes asymptotically good error-correcting codes to extend the basic scheme. We present a simpler construction here which achieves the same efficiency and which (to the best of our knowledge) has not appeared before. Let $G : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ be a pseudo-random generator, where ℓ will be fixed later. The receiver begins by choosing random $r_1, \dots, r_m \in \{0, 1\}^\ell$ and transmitting these to the sender. The sender chooses a random $s \in \{0, 1\}^n$ and responds with $(\bigoplus_{i:M_i=1} r_i) \oplus G(s)$ (where M_i is the i^{th} bit of M). As in the basic Naor scheme, hiding follows easily from the pseudorandomness of G . As for binding, we have

$$\begin{aligned}
& \Pr_{r_1, \dots, r_m} \left[\begin{array}{l} \exists M \neq M', s, s' \text{ such that} \\ (\bigoplus_{i:M_i=1} r_i) \oplus G(s) = (\bigoplus_{i:M'_i=1} r_i) \oplus G(s') \end{array} \right] \\
&= \Pr_{r_1, \dots, r_m} \left[\begin{array}{l} \exists M \neq M', s, s' \text{ such that} \\ \bigoplus_{i:M_i \oplus M'_i=1} r_i = G(s) \oplus G(s') \end{array} \right] \\
&= \Pr_{r_1, \dots, r_m} \left[\begin{array}{l} \exists N \neq 0^m, s, s' \text{ such that} \\ \bigoplus_{i:N_i=1} r_i = G(s) \oplus G(s') \end{array} \right] \\
&\leq \sum_{\substack{s, s' \\ N \neq 0^m}} \Pr_{r_1, \dots, r_m} \left[\bigoplus_{i:N_i=1} r_i = G(s) \oplus G(s') \right] \\
&< 2^m \cdot 2^{2n} \cdot 2^{-\ell}.
\end{aligned}$$

Setting $\ell = n + m + k$, we obtain a scheme that is binding except with probability $\varepsilon_b < 2^{n-k}$ (as previously) and which requires only $O((m+k)/\log S)$ invocations of an S -hard permutation π .

References

- [Blu83] M. Blum. Coin Flipping by Telephone — A Protocol for Solving Impossible Problems. *ACM SIGACT News* 15(1): 23–27, 1983.
- [BM84] M. Blum and S. Micali. How to Generate Cryptographically Strong Sequences of Pseudo-Random Bits. *SIAM J. Computing* 13(4): 850–864, 1984.
- [CW79] J. Carter and M. Wegman. Universal Classes of Hash Functions. *Journal of Computer and System Sciences* 18(2): 143–154, 1979.
- [DDN00] D. Dolev, C. Dwork, and M. Naor. Non-Malleable Cryptography. *SIAM J. Computing* 30(2): 391–437, 2000.
- [Fis02] M. Fischlin. On the Impossibility of Constructing Non-Interactive Statistically-Secret Protocols From any Trapdoor One-Way Function. *Cryptographers' Track — RSA 2002*, LNCS vol. 2271, pp. 79–95, 2002.

- [GGKT05] R. Gennaro, Y. Gertner, J. Katz, and L. Trevisan. Bounds on the Efficiency of Generic Cryptographic Constructions. *SIAM J. Computing* 35(1): 217–246, 2005.
- [G⁺00] Y. Gertner, S. Kannan, T. Malkin, O. Reingold, and M. Viswanathan. The Relationship between Public-Key Encryption and Oblivious Transfer. *41st IEEE Symposium on Foundations of Computer Science (FOCS)*, IEEE, pp. 325–335, 2000.
- [GMR01] Y. Gertner, T. Malkin, and O. Reingold. On the Impossibility of Basing Trapdoor Functions on Trapdoor Predicates. *42nd IEEE Symposium on Foundations of Computer Science (FOCS)*, IEEE, pp. 126–135, 2001.
- [Gol01] O. Goldreich. *Foundations of Cryptography, vol. 1: Basic Tools*. Cambridge University Press, 2001.
- [GGM85] O. Goldreich, S. Goldwasser, and S. Micali. On the Cryptographic Applications of Random Functions. *Advances in Cryptology — Crypto '84*, LNCS vol. 263, Springer-Verlag, pp. 276–288, 1985.
- [GGM86] O. Goldreich, S. Goldwasser, and S. Micali. How to Construct Random Functions. *Journal of the ACM* 33(4): 792–807, 1986.
- [GL89] O. Goldreich and L. Levin. Hard-Core Predicates for any One-Way Function. *21st ACM Symposium on Theory of Computing (STOC)*, ACM, pp. 25–32, 1989.
- [GM84] S. Goldwasser and S. Micali. Probabilistic Encryption. *J. Computer and System Sciences* 28(2): 270–299, 1984.
- [HILL99] J. Hastad, R. Impagliazzo, L. Levin, and M. Luby. A Pseudorandom Generator From any One-Way Function. *SIAM J. Computing* 28(4): 1364–1396, 1999.
- [IL89] R. Impagliazzo and S. Luby. One-Way Functions are Essential for Complexity-Based Cryptography. *30th IEEE Symposium on Foundations of Computer Science (FOCS)*, IEEE, pp. 230–235, 1989.
- [IR89] R. Impagliazzo and S. Rudich. Limits on the Provable Consequences of One-Way Permutations. *21st ACM Symposium on Theory of Computing (STOC)*, ACM, pp. 44–61, 1989.
- [KST99] J.H. Kim, D.R. Simon, and P. Tetali. Limits on the Efficiency of One-Way Permutation-Based Hash Functions. *40th IEEE Symposium on Foundations of Computer Science (FOCS)*, IEEE, pp. 535–542, 1999.
- [Naor91] M. Naor. Bit Commitment Using Pseudorandomness. *J. Cryptology* 4(2): 151–158, 1991.

- [NY89] M. Naor and M. Yung. Universal One-Way Hash Functions and their Cryptographic Applications. *21st ACM Symposium on Theory of Computing (STOC)*, ACM, pp. 33–43, 1989.
- [RTV04] O. Reingold, L. Trevisan, and S. Vadhan. Notions of Reducibility Between Cryptographic Primitives. *1st Theory of Cryptography Conference*, LNCS vol. 2951, Springer-Verlag, pp. 1–20, 2004.
- [Rom90] J. Rompel. One-Way Functions are Necessary and Sufficient for Secure Signatures. *22nd ACM Symposium on Theory of Computing (STOC)*, ACM, pp. 387–394, 1990.
- [Rud88] S. Rudich. Limits on the Provable Consequences of One-Way Functions. Ph.D. thesis, University of California at Berkeley, 1988.
- [Rud91] S. Rudich. The Use of Interaction in Public Cryptosystems. *Adv. in Cryptology — Crypto '91*, LNCS vol. 576, Springer-Verlag, pp. 242–251, 1992.
- [Sim98] D.R. Simon. Finding Collisions on a One-Way Street: Can Secure Hash Functions be Based on General Assumptions? *Adv. in Cryptology — Eurocrypt '98*, LNCS vol. 1403, Springer-Verlag, pp. 334–345, 1998.
- [Yao82] A.C.-C. Yao. Theory and Application of Trapdoor Functions. *23rd IEEE Symposium on Foundations of Computer Science (FOCS)*, IEEE, pp. 80–91, 1982.